What is claimed is:

1. A method for positioning data in a database, comprising:

constructing a base graph in which a plurality of data blocks in the database correspond to respective ones of a plurality of vertices, and in which at least one affinity between objects corresponds to at least one edge; and

constructing a simpler graph, comprising:

a plurality of final vertices, each corresponding to at least two of the vertices from the base graph; and

at least one final edge corresponding to at least one of the edges in the base graph; and

selecting a simple partition for the simpler graph; and

determining a final partition for the data in the database from the simple partition of the simpler graph.

2. The method of claim 1, wherein the base graph and simpler graph are weighted graphs, and wherein the weight of a given edge corresponds to the affinities of the data blocks corresponding to the vertices adjacent the given edge.

3. The method of claim 1, wherein the simpler graph is constructed by constructing one or more intermediate graphs, each intermediate graph being a subsequent graph to one from which it is constructed and being a prior graph to a graph constructed from it, each intermediate graph comprising:

a plurality of new vertices corresponding to pairs of vertices from the previous graph; and

at least one new edge corresponding to at least one edge adjacent to one of the vertices in the pair of vertices corresponding to the new vertex adjacent to the at least one new edge; and

wherein the simpler graph is constructed from one of the one or more intermediate graphs.

4. The method of claim 3, wherein the base graph and simpler graph are weighted graphs, and wherein the weight of an edge corresponds to the affinities of the data blocks corresponding to the vertices adjacent the given edge.

5. The method of claim 1, wherein the determining a final partition for the data in the database from the simple partition of the simpler graph comprises:

determining a penultimate partition for the base graph wherein each group comprises every parent vertex to any daughter vertex in the simpler graph if the group comprises at least one parent vertex to that daughter vertex; and

determining a rough distribution of the data blocks in which a given data block is positioned on a page with each data block with which the given data block shares an adjacent edge, unless that edge is cut in the penultimate partition; and

refining the rough distribution by moving at least one data block from the page on which it was positioned during the determining a rough distribution of the data blocks.

6. The method of claim 1, wherein every data block of data in the database corresponds to a vertex in the base graph.

7. The method of claim 1, wherein fewer than every data block of data in the database correspond to a vertex in the base graph.

8. A method for positioning data in a database, comprising:

constructing a base graph in which a plurality of data blocks in the database each correspond to one of a plurality of vertices, wherein the base graph is a weighted graph having weights of edges corresponding to affinities of data blocks of data; and

constructing a simpler graph, comprising:

a plurality of final vertices, each corresponding to at least two of the vertices from the base graph; and

at least one final edge corresponding to at least one of the edges in the base graph; and

wherein the simpler graph is a weighted graph having edges with weights corresponding to weights of edges in the base graph;

selecting a simple partition for the simpler graph; and

determining a penultimate partition for the base graph wherein each group comprises every parent vertex to any daughter vertex in the simpler graph if the group comprises at least one parent vertex to that daughter vertex; and

determining a rough distribution of the data blocks in which a given data block is positioned on the a page with each data block with which the given data block shares an adjacent edge, unless that edge is cut in the penultimate partition; and

refining the rough distribution by moving at least one data block from the page on which it was positioned during the determining a rough distribution of the data blocks.

9. The method of claim 8, wherein every data block of data in the database corresponds to a vertex in the base graph.

10. The method of claim 8, wherein fewer than every data block of data in the database correspond to a vertex in the base graph.

11. A method of positioning a first new data block of data on a page, comprising:

buffering the first new data block in server memory; and

selecting a page containing data blocks having a high collective affinity for the first new data block; and

determining whether the collective affinity of the data blocks on the selected page for the first new data block exceeds a preselected value; and

positioning the first new data block on the selected page if the collective affinity of the data blocks on the selected page exceed the preselected value; and

positioning the first new data block on a new page if the collective affinity of the data blocks on the selected page do not exceed the preselected value.

12. The method of claim 11, wherein additional new data blocks created at the same time as the first new data block are positioned on the same page as the first new data block.

13. A method for assigning a weight to an edge of a weighted graph corresponding to data blocks in a database, the method comprising using information from the database about explicit connections between at least one pair of data blocks, the at least one pair of data blocks corresponding to at least one pair of vertices of the weighted graph.

14. The method of claim 13, wherein the information comprises extents, relationships, and collections defined by a schema of the database.

15. The method of claim 13, wherein the information is defined by a traversal algorithm.

16. A method for assigning a weight to an edge of a weighted graph corresponding to data blocks in a database, the method comprising using statistical information regarding the patterns of past accesses of data blocks.

17. The method of claim 16, wherein the statistical information is used to select an access type from a pre-defined set, the access type defining the weight.

18. A method for assigning a weight to an edge of a weighted graph corresponding to data blocks in a database, the method comprising using information provided by an application developer.

19. The method of claim 18, further comprising:

providing a schema of the database defining a pre-assigned access type for at least one collection; and

wherein the using information provided by an application developer comprises assigning the at least one weight to the at least one edge corresponding to at least one relationship between data blocks comprising the collection, the at least one relationship being defined by the pre-assigned access type.

20. The method of claim 18, wherein the using information provided by an application developer comprises:

selecting at least one derived access type from a predefined set for each of the collections of the schema of the database; and

assigning at least one weight to at least one edge corresponding to at least one relationship between data blocks comprising the collection, the at least one relationship being defined by the derived access type.

21. The method of claim 20, wherein the selecting at least one derived access type comprises:

collecting statistical data on access patterns of a plurality of data blocks;

selecting from the predefined set the derived access type most closely matching the statistical data.

22. A method for assigning at least one weight to at least one edge of a weighted graph corresponding to data blocks in a database, the method comprising:

providing a schema of the database defining at least one collection; and

selecting at least one derived access type from a predefined set for the at least one collection of the schema of the database; and

assigning at least one weight to at least one edge corresponding to at least one relationship between data blocks comprising the collection, the at least one relationship being defined by the derived access type.

23. The method of claim 22, wherein the selecting at least one derived access type comprises:

collecting statistical data on access patterns of a plurality of data blocks;

selecting from the predefined set the derived access type most closely matching the statistical data.

24. The method of claim 23, wherein the at least one weight is assigned to at least one edge which is adjacent only to vertices corresponding to data blocks having access patterns for which no data was collected.

25. A method for positioning data in a database, comprising:

constructing a weighted base graph in which a plurality of data blocks in the database each correspond to one of a plurality of vertices, and in which at least one affinity between objects corresponds to at least one edge; and

constructing a weighted simpler graph, comprising:

a plurality of final vertices, each corresponding to at least two of the vertices from the base graph; and

at least one final edge corresponding to at least one of the edges in the base graph; and

selecting a simple partition for the simpler graph; and

determining a final partition for the data in the database from the simple partition of the simpler graph; and

wherein at least one weight of at least one edge is assigned using information from the database about explicit connections between at least one pair of data blocks, the at least one pair of data blocks corresponding to at least one pair of vertices of the weighted base graph.

26. The method of claim 25, wherein the information comprises extents, relationships, and collections defined by a schema of the database.

27. The method of claim 25, wherein the information is defined by a traversal algorithm.

28. A method for positioning data in a database, comprising:

constructing a weighted base graph in which a plurality of data blocks in the database each correspond to one of a plurality of vertices, and in which at least one affinity between objects corresponds to at least one edge; and

constructing a weighted simpler graph, comprising:

a plurality of final vertices, each corresponding to at least two of the vertices from the base graph; and

at least one final edge corresponding to at least one of the edges in the base graph; and

selecting a simple partition for the simpler graph; and

determining a final partition for the data in the database from the simple partition of the simpler graph; and

wherein at least one weight of at least one edge is assigned using statistical information regarding the patterns of past accesses of data blocks.

29. The method of claim 28, wherein the statistical information is used to select an access type from a pre-defined set, the access type defining the weight.

30. A method for positioning data in a database, comprising:

constructing a weighted base graph in which a plurality of data blocks in the database each correspond to one of a plurality of vertices, and in which at least one affinity between objects corresponds to at least one edge; and

constructing a weighted simpler graph, comprising:

a plurality of final vertices, each corresponding to at least two of the vertices from the base graph; and

at least one final edge corresponding to at least one of the edges in the base graph; and

selecting a simple partition for the simpler graph; and

determining a final partition for the data in the database from the simple partition of the simpler graph; and

wherein at least one weight of at least one edge is assigned using information provided by an application developer.

31. The method of claim 30, further comprising:

providing a schema of the database defining a pre-assigned access type for at least one collection; and

assigning the at least one weight to the at least one edge corresponding to at least one relationship between data blocks comprising the collection, the at least one relationship being defined by the pre-assigned access type.

32. The method of claim 30, further comprising:

selecting at least one derived access type from a predefined set; and

assigning the at least one weight to the at least one edge corresponding to at least one relationship between data blocks comprising the collection, the at least one relationship being defined by the derived access type.

33. The method of claim 32, wherein the selecting at least one derived access type comprises:

collecting statistical data on access patterns of a plurality of data blocks;

selecting from the predefined set the derived access type most closely matching the statistical data.

34. A method for positioning data in a database, comprising:

providing a schema of the database defining at least one collection and at least one predefined set of derived access types; and

constructing a weighted base graph in which a plurality of data blocks in the database each correspond to one of a plurality of vertices, and in which at least one affinity between objects corresponds to at least one edge, the constructing comprising:

selecting at least one derived access type from the predefined set for the at least one collection of the schema of the database; and

assigning at least one weight to at least one edge corresponding to at least one relationship between data blocks comprising the collection, the at least one relationship being defined by the derived access type; and

constructing a weighted simpler graph, comprising:

a plurality of final vertices, each corresponding to at least two of the vertices from the base graph; and

at least one final edge corresponding to at least one of the edges in the base graph; and

selecting a simple partition for the simpler graph; and

determining a final partition for the data in the database from the simple partition of the simpler graph.

35. The method of claim 34, wherein the selecting at least one derived access type comprises:

collecting statistical data on access patterns of a plurality of data blocks; and

selecting from the predefined set the derived access type most closely matching the statistical data.

36. The method of claim 35, wherein the at least one weight is assigned to at least one edge which is adjacent only to vertices corresponding to data blocks having access patterns for which no data was collected.

37. The method of claim 35, wherein the simpler graph is constructed by constructing one or more intermediate graphs, each intermediate graph being a subsequent graph to one from which it is constructed and being a prior graph to a graph constructed from it, each intermediate graph comprising:

a plurality of new vertices corresponding to pairs of vertices from the previous graph; and

at least one new edge corresponding to at least one edge adjacent to one of the vertices in the pair of vertices corresponding to the new vertex adjacent to the at least one new edge; and

wherein the simpler graph is constructed from one of the one or more intermediate graphs.

38. The method of claim 35, wherein the determining a final partition for the data in the database from the simple partition of the simpler graph comprises:

determining a penultimate partition for the base graph wherein each group comprises every parent vertex to any daughter vertex in the simpler graph if the group comprises at least one parent vertex to that daughter vertex; and

determining a rough distribution of the data blocks in which a given data block is positioned on the a page with each data block with which the given data block shares an adjacent edge, unless that edge is cut in the penultimate partition; and

refining the rough distribution by moving at least one data block from the page on which it was positioned during the determining a rough distribution of the data blocks.

39. A method for positioning data in a database, comprising:

providing a schema of the database defining a first and second collection and at least one predefined set of derived access types; and

constructing a weighted base graph in which a plurality of data blocks in the database each correspond to one of a plurality of vertices, and in which at least one affinity between objects corresponds to at least one edge, the constructing comprising:

collecting statistical data on access patterns of a plurality of data blocks associated with the first collection; and

selecting from the predefined set the derived access type most closely matching the statistical data; and

assigning at least one weight to at least one edge corresponding to at least one relationship between data blocks associated with the second collection, the at least one relationship being defined by the derived access type; and

constructing one or more intermediate graphs, each intermediate graph being a subsequent graph to one from which it is constructed and being a prior graph to a graph constructed from it, each intermediate graph comprising:

a plurality of new vertices corresponding to pairs of vertices from the previous graph; and

at least one new edge corresponding to at least one edge adjacent to one of the vertices in the pair of vertices corresponding to the new vertex adjacent to the at least one new edge; and

constructing a weighted simpler graph from one of the intermediate graphs, the simpler graph comprising:

a plurality of final vertices, each corresponding to at least one of the vertices from the base graph; and

at least one final edge corresponding to at least one of the edges in the base graph; and

selecting a simple partition for the simpler graph; and

determining a penultimate partition for the base graph wherein each group comprises every parent vertex to any daughter vertex in the simpler graph if the group comprises at least one parent vertex to that daughter vertex; and

determining a rough distribution of the data blocks in which a given data block is positioned on the a page with each data block with which the given data block shares an adjacent edge, unless that edge is cut in the penultimate partition; and

refining the rough distribution by moving at least one data block from the page on which it was positioned during the determining a rough distribution of the data blocks.

40. The method of claim 39, further comprising:

buffering a first new data block in server memory; and

selecting a page containing data blocks having a high collective affinity for the first new data block; and

determining whether the collective affinity of the data blocks on the selected page for the first new data block exceeds a preselected value; and

positioning the first new data block on the selected page if the collective affinity of the data blocks on the selected page exceed the preselected value; and

positioning the first new data block on a new page if the collective affinity of the data blocks on the selected page do not exceed the preselected value.

41. The method of claim 40, wherein additional new data blocks created at the same time as the first new data block are positioned on the same page as the first new data block.

42. A method of positioning a first new data block of data on a page, comprising:

buffering the first new data block in server memory;

selecting a page containing data blocks having a high collective affinity for the first new data block; and

positioning the first new data block the page containing data blocks having a high collective affinity for the first new data block.

43. The method of claim 42, wherein additional new data blocks created at the same time as the first new data block are assigned relatively high affinities for one another.